

AD-A114 119

BELL NORTHERN RESEARCH LTD OTTAWA (ONTARIO)

NETWORK SPECIFICATIONS FOR A VIRTUAL SPACE TELECONFERENCING SYS--ETC(11)


APR 82

F/G 17/2


MDA903-81-C-0180

NL

UNCLASSIFIED




END
DATE
FILMED
6-82
DTIC



12

AD A114119



DTIC FILE COPY

DTIC
ELECTE
MAY 5 1982
S A D

Prepared by
Bell-Northern Research Ltd.
Ottawa, Ontario

This document has been approved
for public release and sale; its
distribution is unlimited.

82 05 04 088

APRIL 1982

FINAL REPORT

Network Specifications For A Virtual
Space Teleconferencing System

PREPARED FOR:

Defense Advanced Research Projects Agency

Contract No.: MDA903-81-C-0180
BNR Ref. No.: P379/TR6735

PREPARED BY:

Bell-Northern Research Ltd.
P.O. Box 3511, Station C
Ottawa, Ont.
K1Y 4H7

RECEIVED
APR 16 1982
A

ALL INFORMATION CONTAINED
HEREIN IS UNCLASSIFIED
DATE 01-11-81 BY 1045

Table of Contents

1. INTRODUCTION	1
2. THE CONFERENCING ENVIRONMENT	2
2.1 THE VIDEO SPACE	2
2.2 THE AUDIO SPACE	4
2.3 THE GRAPHICS SPACE	5
3. THE NETWORK REQUIREMENTS	6
4. SOME CANONICAL NETWORKS	11
4.1 TERRESTRIAL NETWORK DESIGN	11
4.1.1 TERRESTRIAL NETWORK DESIGN : ONE SITE PER CLUSTER .	11
4.1.2 TERRESTRIAL NETWORK DESIGN :	
MORE THAN ONE SITE PER CLUSTER	17
4.1.3 THE SUBCONFERENCING PROBLEM :	
MORE THAN ONE SITE PER CLUSTER	18
4.2 SATELLITE NETWORK DESIGN	18
4.3 FLEXIBILITY COMPARISONS FOR THE NETWORKS	19
5. THE INTRA CLUSTER NETWORK	22
6. THE SITE CONTROLLER	26
7. THE CLUSTER CONTROLLER	27
8. A DEMONSTRATION NETWORK	28
9. CONCLUSIONS	32
10. FUTURE WORK	33
FIGURES	34



Letter on file

A

NETWORK SPECIFICATIONS FOR A VIRTUAL SPACE TELECONFERENCING SYSTEM

1. INTRODUCTION

This note outlines the network requirements to provide a flexible experimentation environment for a special video/visual teleconferencing system. To make this as selfcontained as possible Section 2 will provide a description of the conferencing environment. This will carefully define the concept of a "virtual space", a "pseudo virtual space", and various "partial virtual spaces". Each of these conferencing environments imposes different requirements on the interconnecting nodal network. Section 3 will describe the connectivities, capacities, and switching requirements for each of the conferencing environments. Section 4 will introduce a set of canonical network topologies and capabilities that can be configured to satisfy the requirements of the various spaces. Section 5 describes in more detail the requirements and a potential design for a local network to support a cluster of teleconferencing nodes. Section 6 describes the internodal network. Section 7 describes the cluster controller for the system. Finally Section 8 describes the requirements for a demonstration network. Section 9 concludes and Section suggests future research.

The Figures are all at the end of the report.

2. THE CONFERENCING ENVIRONMENT

The ground rules for the conference system are as follows:

- * There is one participant at each conference site.
- * There are N conference sites. For the initial design $N=5$.
- * There are three modalities in the conference.
 - a "video space" affording views of the participants
 - an "audio space" affording voice distribution to all
 - a "graphics space" for the common display of documents

The various spaces come in different shapes. The next sections discuss these variations.

2.1 THE VIDEO SPACE

In order to describe the video space it is useful to imagine a model of 5 participants sitting around a round table as in Figure 2.1:1. Each participant may look around the room to see the fellow conferees and while looking at them be able to determine who is looking at whom. Furthermore, because of the geometry, participant 1 will see images different than any other participant. As can be seen, by shifting the eyes slightly there is a new view. In fact there are an infinite number of views for each participant. However it is probably true that the views involving just one other participant are important. Thus $V(1,4)$ is the view seen by participant 1 when looking at participant 4. Evidently when the system is discretized in this way there are $N(N-1)$ or 30 ($N=5$) distinct views. A teleconferencing system involving participants at geographically separated sites that provides all $N(N-1)$ views is said to be a "virtual space". A teleconferencing system may provide from 1 (e.g. voice-switched video) to a full set of views.

When the live conferencing situation is replaced by the remote teleconferencing system, two very important video aspects need to be evaluated. The first is the quality of the pictures of the remote participants that are displayed and the second is the "presence" of those participants. The picture quality is primarily determined by the bandwidth available per view, and the ability of the video coder to exploit that bandwidth. The "presence" of the remote participants is primarily determined by the individual's ability to establish eye contact and the realism of the views. Since the network has a limited capacity, there is a tradeoff between capacity per view and the total number of views provided.

Figure 2.1:2 shows a (tele)conferencing site that would be the first step in providing a remote implementation to the live model above. In this figure participant 1 is the only live conferee at the site. The other participants have been replaced by monitors, to show their image, and cameras to act as their eyes. Note that the cameras are all pointed towards participant 1 as would be the eyes of the other participants

while looking at 1. A full conference would involve 5 sites identical to the one in Figure 2.1:2.

The amount of presence in the conferencing system is dependent on the number of views provided. Thus the full presence system is defined as follows:

- * A "Virtual Space" is a teleconferencing system that provides view $V(i,j)$ to site i on monitor j for all i not equal to j . Note that all sites transmit and receive exactly $N-1$ views and that the network carries $N(N-1)$ simultaneous views.

A "Pseudo Virtual Space" can be created by noting that a participant is able to look only at one participant at a time and hence need only be provided with that view and not any of the others. Thus if 1 is looking at 4, then only view $V(1,4)$ need be provided. This means that when 1 looks at 4, the output of the camera at site 4 in position 1 must be sent to site 1. Thus the action of looking at a new participant at site 1 triggers the sending (after some delay), of the appropriate view from that remote site. Thus

- * A "Pseudo Virtual Space" is a teleconferencing system that provides for every site i , the view $V(i,j)$ if participant i is looking at participant j . Note that this means that any site receives exactly one view, that any site may transmit up to $N-1$ views, and that the network must carry N simultaneous views. The network must be able to switch views.

The major advantage of both the Virtual and Pseudo Virtual Spaces are that they are able to provide eye contact and indications of who is looking at whom. It is possible to give up the eye contact but retain information about who is looking at whom. This would be done using the arrangement in Figure 2.1:2 but transmit $V(j,i)$ to all participant sites if participant i is looking at j . Note that $V(j,i)$ originates at site i . The information about who is looking at whom is relayed digitally, and displayed graphically, to all sites. This will be defined as the "Who/Whom Virtual Space". Thus

- * The "Who/Whom Video Space" is a teleconferencing system that provides for every site i , the view $V(j,i)$ to all other sites when i is looking at j . Note that every site transmits exactly one view and every site receives $N-1$ views. The network must carry exactly N simultaneous views but must broadcast them to all sites.

As before a "Pseudo Who/Whom Video Space" is defined by noting that a participant can only be looking at one place at a time. Thus if i is looking at j , then i will transmit $V(j,i)$ and i will receive the view $V(k,j)$ that is being transmitted by site j . Note that a view need be transmitted from a site unless there is at least one participant looking at that site. Thus

- * The "Pseudo Who/Whom Video Space" is a teleconferencing system that provides the view $V(j,i)$ from site i to any other site looking at i when participant i is looking at j . Note that every site receives exactly one view and will transmit a view if there is at least one other participant looking there. The network must be able to carry up to N simultaneous views. It will carry, depending on the view patterns, from 2 to N simultaneous views.

If both the eye contact and the who/whom aspect are given up the arrangement in Figure 2.1:3 is sufficient. In this format the the video image of all the participants is available at all times. There is no indication as to who is looking at whom but there is pseudo eye contact. The system employs only a single camera so that each site is capable of only a single view, $V(*,i)$ which is transmitted to all other sites. Thus

- * The "Continuously Available Video Space" is a teleconferencing system that provides the view $V(*,i)$ to all other sites from site i . Note that each site receives $N-1$ views and transmits 1 view. The network must carry N simultaneous views and broadcast them to each site.

For completeness, a final possibility should be mentioned. In this system even the continuous availability is given up and only one view is transmitted. In a teleconferencing mode this view usually originates with the speaker and is implemented using voice switching. This is defined as a "Voice Switched Video Space".

2.2 THE AUDIO SPACE

The variations on the audio space are much less than those in the video space. Each site has only one audio source and that must be transmitted to all other sites. A "Simple Full Audio Space" implies that the voices can be coordinated with the views. For the systems shown in Figures 2.1:1 and 2.1:2, this means that the voice streams must be kept separate and fed to speakers associated with each monitor. For the case in Figure 2.1:3, it may be possible to obtain the separation using stereo speakers. However the mixing ratios are site dependent and would require complex mixing, especially when sophisticated voice coding

is used. For the present it is assumed that this type of mixing will not be used.

For the "Simple Full Audio Space", each site will transmit one voice signal and will receive $(N-1)$ voice signals. The network will thus be required to broadcast N simultaneous voice signals. This system assumes that all of the voice channels are of identical, presumably high, quality. However normal meeting protocol suggests that only one person at time, except for interruptions, will be talking. This suggests an "Enhanced Full Audio Space" with one switched high quality voice channel for the speaker and $(N-1)$ lower quality channels for the interrupters. In this case all but the speaker's site will transmit one low quality channel, and receive $(N-2)$ low quality channels along with one high quality channel. The speaker site will transmit one high quality channel, and receive $(N-1)$ low quality channels. The network must simultaneously carry and broadcast one high quality channel, and $(N-1)$ low quality channels. If the high quality channel merely augments the low quality channels, then the network can broadcast, without switching, the N low quality channels. However, if the high quality channel is an incompatible coding scheme, then both the low and high quality channels must be switched.

2.3 THE GRAPHICS SPACE

The "Graphics Space" is meant to distribute documents, and facilitate discussions about shared images. There is evidently a single input source from each site and all the inputs must be broadcast to all other sites. Thus each site must transmit its own graphics stream and receive $(N-1)$ graphics streams. Since the graphics streams are to be interpreted as a single composite stream at each site, the network must provide some coordination. The network must carry N simultaneous graphics streams (or one composite stream).

3. THE NETWORK REQUIREMENTS

In Section 2 a range of video and audio spaces that could be used as the basis of a teleconferencing environment was given. The list is as follows:

Video Space

Virtual Space	(VS)
Pseudo Virtual Space	(PVS)
Who/Whom Video Space	(WVS)
Pseudo Who/Whom Video Space	(PWVS)
Continuously Available Video Space	(CAVS)
Voice-Switched Video Space	(VSVS)

Audio Space

Full Audio Space	(FAS)
Enhanced Audio Space	(EAS)

Graphics Space

(GS)

The derivation of the network carrying capacity for the range of video and audio spaces above is easily derived from the requirements on distributing views from each site. For instance, VS requires that each site transmit $N-1$ views, a different one to each of the remote sites. In addition, it receives $N-1$ views, one from each site. This results in $N-1$ outward 2-point simplex connections and $N-1$ inward 2-point simplex connections. Because of the site connection symmetry, it is possible to realize these $2*(N-1)$ simplex connections as $N-1$ 2-point full duplex connections. A second example would be PVS. In this case, since the participant can be looking only in one place, there is exactly 1 inward 2-point simplex connection. Since this may originate at any site the input end of the 2-point simplex connection must be switchable amongst any of the other $N-1$ remote sites. Since up to $N-1$ of the other participants may be looking at any single site, there is a need for the potential of $N-1$ outward 2-point simplex connections. Because of the unbalance between inward and outward requirements, the potential for $N-1$ 2-point full duplex connections will be required. A third example is the full audio space (FAS). There is exactly one voice channel for each site but that channel must be distributed to all sites. This is defined as a $(1, N-1):a$ simplex connection, i.e. a connection that connects "1" point to " $N-1$ " points and has a bandwidth " a " kbps. Thus if the video coder requires v kbps (for any of the video spaces) then a simplex 2-point connection of bandwidth " v " is $(1, 1):v$. A 2-point full duplex connection is $(1, 1):v;2$. Using this notation it is quite easy to state the site requirements for each conferencing form. This is shown in Table 3.1

TABLE 3.1

Site Connection Requirements

	Simplex				Full Duplex
	Inward Connections		Outward Connections		
Video Spaces					
VS	N-1	(1,1):v	N-1	(1,1):v	N-1 (1,1):v;2
PVS	1	(1,1):v	0 to N-1	(1,1):v	N-1 (1,1):v;2
WVS	N-1	(1,N-1):v	1	(1,N-1):v	undefined
PWVS	1	(1,1):v	1	(1,k):v, 0 < k < N	undefined
CAVS	N-1	(1,N-1):v	1	(1,N-1):v	undefined
VSVS	1	(1,N-1):v	0 to 1	(1,N-1):v	undefined
Audio Spaces					
FAS	N-1	(1,N-1):a	1	(1,N-1):a	undefined
EAS	1	(1,N-1):b	1	(1,N-1):b	undefined
	+	1 (1,N-1):a	0 to 1	(1,N-1):a	undefined
Graphics Space	N-1	(1,N-1):g	1	(1,N-1):g	undefined

A point worth noting about Table 3.1 is that the number of full duplex connections are listed as "undefined" for anything other than (1,1) types. The reason for this is that although there is a perfectly well defined notion for a broadcast connection (1,N), there is no universal notion for a collection connection (N,1) which would be the return connection in the full duplex case. In fact a collection connection only makes sense if there is associated with it a contention resolution mechanism. There are several that could be used including polling, random access, or global queues. Along with the connections there will be a need to switch either one or both ends of the connection. There are three types of switching that need to be distinguished. To make the distinction it is necessary to introduce the idea of a "port". A view is presented to the network at an output port. A view is absorbed from the network at an input port. Figure 3.1:1 shows two sites each with two input ports and two output ports. If it were possible to get the sites together so that the output port of site A could be mated with the input port of site B as shown in Figure 3.1:2, then a view could be exchanged over the mated or connected ports. However, this is not possible. Consequently a simplex connection consisting of one input port attached to an output port is used to create the remote mating as shown in Figure 3.1:3.

Using these notions the three types of switching are defined as follows:

Local Selection:

The input port of the connection is switched amongst the output ports of the site or the output port of a connection is switched amongst the input ports of the site. This is the most benign type of switching.

Locally Induced Switching:

The local input port of the connection is switched to the output port of another site or the local output port of the connection is switched to the input port of another site. This does not require any cross network connection to induce the switch.

Remotely Induced Switching:

The remote input port of the connection is switched to the output port of another site or the remote output port of the connection is switched to the input port of another site. This requires a cross network connection to induce the switch.

The three types of switching have increasing delay and complexity. It is easy to see that VS has neither switching nor selection and that PVS has remotely induced switching. WVS has local selection, to determine what view to send to the remote users while PWVS has both local selection and remotely induced switching. It is also easy to see that VSVS has only remotely induced switching. The summary of the switching requirements for the various cases is given in Table 3.2.

TABLE 3.2

Switching Requirements

	Local Selection	Locally Induced Switching	Remotely Induced Switching
Video Spaces			
VS	No	No	No
PVS	No	No	Yes
WVS	Yes	No	No
PWVS*	Yes	No	Yes
CAVS	No	No	No
VSVS	No	No	Yes
Audio Spaces			
FAS	No	No	No
EAS	Yes	No	Yes
Graphics Space	No	No	No

An interesting point in Table 3.2 is that there is no locally induced switching. This means that all switching, other than local selection, will require cross network signalling to effect the change.

Based on the site connection requirements of Table 3.1 and noting the symmetry of the connections required to transport the views it is easy to derive the network requirements for the various cases. These are summarized in Table 3.3

TABLE 3.3

Network Connection Requirements

	Simplex Connections	Full Duplex Connections
Video Spaces		
VS	$N(N-1)$ (1,1):v	$N(N-1)/2$ (1,1):v;2
PVS	N (1,1):v	N (1,1):v;2
WVS	N (1,N-1):v	undefined
PWVS*	N (1,1):v	N (1,1):v;2
CAVS	N (1,N-1):v	undefined
VSVS	1 (1,N-1):v	undefined
Audio Spaces		
FAS	N (1,N-1):a	undefined
EAS	! N (1,N-1):b	undefined
	+ 1 (1,N-1):a	undefined
Graphics Space	N (1,N-1):g	undefined

* This is a simplified requirement. If site i is being viewed by k remote sites, then a (1,k):v connection is required. Thus as little as 2 and as many as N connections may be required. See Table 3.1

To define an actual teleconferencing environment a choice is made from the video, audio, and graphics spaces. Thus (VS, FAS, GS) and (VS, EAS, GS) are two distinct teleconferencing environments. A network must be designed to satisfy a complete environment. From the table it can be seen that the network and site requirements vary considerably amongst the various alternatives. As will be seen later, a network that is designed to satisfy one of the alternatives does not necessarily, even with the addition of switching, do very well on the the others.

4.

SOME CANONICAL NETWORKS

This section introduces some network configurations that can be used to satisfy the network requirements of the previous section. The major concern of this section is the topological design of the network. A network connection is made up of links. A simplex link is just a fixed unswitched $(1,1):v$ connection between two points. It may take several links to make up a connection. Both terrestrial and satellite networks will be considered. Terrestrial networks are characterized by the obvious fact that the connections required for the various conferencing environments are made up of links that are constrained to connect only 2 physical locations. Thus a $(1,N):v$ connection requires at least $N-1$ $(1,1):v$ or $(1,1):v;2$ links. Satellite networks can be viewed as supplying $(1,k):z$ of capacity "z". It is interesting that the space craft capacity used by a $(1,k):z$ link is only "z" and is independent of k. Furthermore it is usually easy to take the two links $(1,k):z$ and $(1,k):x$ and produce a single link $(1,k):(z+x)$. The immediate implication of this observation is that the bandwidth and connectivity freedom with a satellite network is far greater than that of a terrestrial network.

4.1

TERRESTRIAL NETWORK DESIGN

The most important input for the terrestrial network design is the geographical location of the sites. In order to have a general framework for the design of a very versatile conferencing system the notion of a "local cluster" of sites is introduced. A set of sites is said to comprise a local cluster if they are sufficiently close as to allow analog video transmission between them. This analog video, when combined with the voice and the graphics space data will make up a local network. The design of the local network will be discussed in Section 6. This section will be concerned with the inter cluster network design problem.

Section 4.1.1 will discuss the terrestrial inter cluster design problem when there is only one site per cluster. Section 4.1.2 will discuss the general design problem.

4.1.1

TERRESTRIAL NETWORK DESIGN: ONE SITE PER CLUSTER

The basic problem is to discover a collection of links that satisfy the network requirements of Section 3 at the minimal cost. The major lesson of this section will be that a network that satisfies the requirements for, say VS, will not provide any more capacity for PVS even with the addition of unlimited switching.

There will be no attempt to use actual tariffs in this exercise. However the results will make it quite easy to introduce real tariffs. The reason for this is that the network requirements in Section 3 constrain the possibilities sufficiently that it is only necessary to consider a few canonical networks. The major constraining force is the flow symmetry amongst the sites. The one site per cluster case is

important because it maintains that symmetry.

For the network design it will be assumed that either simplex (1,1):z or duplex links (1,1):z;2 are available. However most digital offerings by the carriers are full duplex only. In fact simplex connections appear to be available only on satellites. Half duplex (two way alternate) links will not be considered. Upon examination of the three tables in Section 3 it can be seen that topologically (i.e. ignoring the bandwidth differences) various spaces are equivalent. The equivalences are shown below. These networks will be referred to as the "Service Networks". These service networks embody the flow constraints of the various conferencing spaces.

SERVICE NETWORKS

"Fully Connected (FCN)" : VS
"Dynamic Allocation (DAN)": PVS
"Complete Broadcast (CBN)": WVS, FAS, EAS-2, CAVS, GS
"Switched Complete Broadcast(SCBN)": PWVS
"Switched Broadcast (SBN)": VSVS, EAS-1

The equivalences are arbitrarily named as indicated. The optimal network topologies for the various equivalent networks are now derived.

Fully Connected Network (FCN)

Table 3.2 shows that there are exactly $N(N-1)$ (1,1);x non switched connections. This means that there are at least $N(N-1)$ (1,1);x links. Since it is possible to realise the network with only $N(N-1)$;x links, the network total distance is minimized when the cluster are directly connected. Because of the symmetry, this can be done with $N(N-1)$ (1,1);x links or $N(N-1)/2$ (1,1);x:2 links. This canonical network will be referred to as a "Fully Connected Network".

Dynamic Allocation Network (DAN)

Table 3.2 indicates that there is an imbalance between the inward and outward connections. Because of the following:

- There must be a link for all POTENTIAL connections in the network*.
- The total number of inward and outward connections within the network must match.
- A link can connect only two fixed geographical locations*.
- A simplex link (1,1), can carry data in one direction only between two fixed geographical locations. A full duplex link is two simplex links between the same geographical locations but oppositely oriented.

It is clear that the minimum number of links in the network is at least $N(N-1) (1,1);x$. The network design with simplex links is quite easy when one notes that the minimum distance solution to connecting N nodes is an $N-1$ link minimum spanning tree. Thus any site can receive a signal from any one of the other sites along this spanning tree if the simplex nodes are oriented to lead to it. An example is given in Figure 4.1:1 where the simplex links are oriented so as to allow "collection" at site A. Note that there are 7 sites in this example and there are 6 links. From Table 3.2, it is seen that PVS requires remotely induced switching. It is clear that although the network in Figure 4.1:1 will allow collection of the remote views at A, its unidirectionality prevents A from inducing the changes. This problem will be taken care of later.

To complete the solution, it is only necessary to lay down 6 more identical, except for orientation, minimum spanning trees in Figure 4.1:1 to obtain the network in Figure 4.1:2. The $\{x,y\}$ on the links mean that there are x simplex links in the direction of the arrow and y in the opposite direction. In all cases $x+y=7$. The numbers are easily derived by partitioning the network at the link and counting the number of sites on the side of the network in the direction of the arrow. Furthermore all the links will be occupied on any cut if the sites on one side receive views from the other. Hence, in the worst case there is never any residual capacity in any of the composite links. Thus the network is optimal over all possible networks.

The question remains as to what to do about the requirement of the remotely induced switching. There are two possibilities. The first is

* This statement is, interestingly enough, not true for satellite networks.

to note that the switching information requires very little capacity and that all composite links have at least one simplex link in both directions. Thus it is possible to rob a small amount of the capacity within the tree for signalling. A second possibility results from the observation that if the orientation of the links in Figure 4.1:1 are reversed the network implements a (1,6) "broadcast" connection of the variety needed by the Complete Broadcast Network. Furthermore, using similar arguments to those above, it is clear that this topology is optimal for the Complete Broadcast Network. Since all complete conferencing systems will require a broadcast subnetwork for the voice and graphics spaces, this network could also be used for video space signalling.

This network topology will be referred to as the Simplex Tree Network. The specific version with the links oriented for dynamic allocation will be qualified as Simplex Tree (Collection) Network. The topology of the Simplex Tree will be seen to be optimal for several of the canonical network cases derived above.

If the network is designed using full duplex links, the natural place to start is with the minimum spanning tree solution for simplex links already derived. To realize a composite link of Figure 4.1:2 using full duplex links, it is clear that it will require the maximum of $\{x,y\}$ duplex links to realize the independent simplex links. This gives the link capacities in Figure 4.1:3. Note that the total number of full duplex links is 33 or 66 simplex links. The optimal Simplex Tree required only 42 links. Thus the Duplex Tree network requires much more link capacity than the Simplex Tree. Furthermore, at most cuts there is residual capacity flowing from one side or the other. This suggests that if this excess capacity were used to multiply connect the network partitions that the total number of links would be reduced.

This is indeed the case. In fact it is possible to connect the 7 nodes in Figure 4.2:3 with only 42 simplex or 21 full duplex links. This is a Ring Network and is shown in Figure 4.1:4. This network will provide a total of 6 rings, 3 in each direction. It is easy to show that a minimum distance routing strategy will allow the 7 connections as required in a Dynamic Allocation Network. It is interesting to note that the network in Figure 4.1:4 may NOT be less expensive than that in Figure 4.1:3. The determining factor is the relative costs of the distance portion versus the fixed portion of a tariff.

For the general solution using full duplex links consider the following observations:

- Because of the flow requirements, each node must have at least $N-1$ full duplex links connected to it. This means that the total number of full duplex network links is at least $N(N-1)/2$.

- A "duplex circuit" of k nodes is a set of $k-1$ duplex links connecting all the nodes together in a loop so that it is possible to travel through each node exactly once and return to the origin. This is essentially a ring of k nodes. This will be called a " k -circuit".
- Any network made up of $(N-1)/2$ N -circuits satisfies the constraints that there are exactly $N-1$ full duplex links for each node. This is easily seen since each circuit gives exactly two links per node.
- For N prime, all networks that satisfy the in/out link constraints are made up from N -circuits. For N prime there are exactly $(N-1)/2$ distinct circuits.
- The N -circuit that minimizes the distance amongst N -circuits is the solution to the travelling salesman problem and is sometimes known as the grand tour.
- For N odd, $(N-1)/2$ grand tour N -circuits make up the minimum distance network. This results in $N-1$ complete rings.
- For N even, only $(N-2)/2$ N -circuits are possible. This leaves $N(N-1)/2 - N(N-2)/2 = N/2$ links left over. The only symmetric way of connecting these nodes is to place them between alternate pairs of nodes as shown in Figure 4.1:5. There are evidently two ways of doing this. The optimum one is evidently the one with the lowest cost. This network and the previous N odd ring will be called a Modified Ring Network.
- It is possible to show that a minimum distance routing algorithm will satisfy the connection requirements for Dynamic Allocation.
- The Modified Ring Network is the minimum cost network for those using $N(N-1)/2$ full duplex links.

The major results of this section are the derivation of two canonical networks that are optimal for Dynamic Allocation Networks. The first, using simplex links is the Simplex Tree. This has the added advantage that although the switching is still remotely induced, the implementation of the switching at the remote node involves local selection only. This is a result of the independence of the N subtrees that make up the total network. For full duplex networks, the optimal (amongst $N(N-1)/2$ link networks) is the Modified Ring Network. This network allows signalling in an easier manner than the Simplex Tree but has a slightly more complex routing algorithm. The switching can be accomplished using local selection.

Complete Broadcast Network (CBN)

The Complete Broadcast Network requires that the output of each node be broadcast to all other nodes. Obviously the Simplex Tree with the links oriented out of the transmitting node to the end of the tree will be the optimal network using simplex links. This canonical network will be called a Simplex Tree (Broadcast) Network. For duplex links the optimal $N(N-1)$ link network is the Modified Ring. For N odd, the routing is just the minimum distance to each of the remaining $N-1$ nodes. This results in two $(1, (N-1)/2)$ connections in opposite directions around the ring. Each of these connections uses $(N-1)/2$ "clockwise" links and $(N-1)/2$ "counterclockwise" links. Since there are N nodes, all the $N(N-1)/2$ full duplex links are counted. For N even, the Modified Ring Network is also optimal. The routing is similar to the N odd case except the extremal nodes are distance $N/2$ and $(N/2)-1$ links away. Note also that the modified ring alternates between $N/2$ and $(N/2)-1$ full duplex links. The routing rule is to connect to the node $N/2$ links away in the direction of the $N/2$ composite link, and to the $(N/2)-1$ node in the direction of the $(N/2)-1$ composite link. It requires a little effort but it is possible to show that the links are all used appropriately. An example is given for $N=6$ in Figure 4.1:7. A careful examination of the routing structure indicates that the construction extends to all even N . The two Complete Broadcast examples show that it is possible effect the switching in the dual Dynamic Allocation case, with local selection.

Switched Complete Broadcast Network (SCBN)

From Table 3.1 it is seen that only one inward and outward link is required for each node. This implies that the network requires at least N simplex links. However, full connectivity for a simplex network requires at least one N -circuit using N simplex links. It is clear that at least 2 N -circuits are needed to handle the case where one site is broadcasting to the even numbered sites and one is broadcasting to the odd numbered sites in an N -circuit. Thus it is clear that at least $2N$ simplex or N duplex links are required using N -circuits. In fact it is possible to satisfy the connectivity constraints with $N-1$ full duplex links (or $2*(N-1)$ simplex links) all homing on to a single site with a switch. This is a "Star Network". The optimum place to put the switch is at the "median" of the network. Note that in this case the simplex and duplex networks are identical.

Switched Broadcast Network (SBN)

As seen from the tables there is only one $(1, N-1)$ connection but that it may originate from any of the N sites. Thus each site must have one inward and outward link. Although the network has to carry only one signal at a time the topology is identical to the "Star Network" derived for the Switched Complete Broadcast Network. This identical topology is only true for terrestrial networks. A satellite network requires only 1 unit of capacity in the SBN case but N units of capacity in the SCBN

case.

Network Summary

The results of this section show that, of all the networks that might have been considered, classic topologies for the most part are optimal. The following are the "Canonical Networks":

CANONICAL NETWORKS

Fully Connected Network
Simplex Tree (Collection) Network
Simplex Tree (Broadcast) Network
Duplex Tree Network
Modified Ring Network
Star Network

Unfortunately, no single network topology is optimal for all three spaces within the conferencing service.

4.1.2 TERRESTRIAL NETWORK DESIGN - MORE THAN ONE SITE PER CLUSTER

The general network design problem recognizes that in some cases the sites will be sufficiently close to each other that they will be able to be linked via analog video. In this case the sites within the cluster will be connected via a cable system to a "cluster controller" that will be the interface to the intercluster network. In this section the concern will again be only with the topological design of the intercluster network.

Unfortunately there is little that can be said about the design of this general problem at this time. The key arguments in Section 4.1.1 depended upon the symmetry of the network. This led to the classical networks in that section. The general problem does not have this symmetry. If there are M clusters and $K(i)$ sites per cluster for a total of N sites, then the arguments in the previous section on site connection requirements and network connections still hold with respect to N . A network that just deletes the links between the $K(i)$ sites at cluster i and leaves the intercluster composite links will satisfy the conditions. However, it is easy to show that there exist better networks that will satisfy the conditions.

The conjecture at this time on the optimal network is that it will consist of a combination of tree and rings. For small networks, the number of options are sufficiently small that they can be enumerated. This means that a solution will be possible.

4.1.3 THE SUBCONFERENCING PROBLEM: MORE THAN ONE NODE PER CLUSTER

For completeness, the general conferencing problem will now be described. This consists of M clusters, each with $K(i)$ sites per cluster. The total number of sites is N . The participant level per site is defined as P . For the site in Figure 2.1:1, $P=5$. All the cases that have been considered so far have had $P=N$. The general problem has $P < N$. There are several levels of service for a network in this situation.

- The network will allow 1 conference of up to P sites.
- The network will allow up to S simultaneous conferences, each of size $< P$, with the sites in each conference being mutually exclusive. S is obviously $< N/P$. A different network will result for each value of S .

The latter situation is the most general. There are no results as yet on this case.

4.2 SATELLITE NETWORK DESIGN

There are basically two cases for the satellite network design. The first occurs when the cluster controller is co-located with the earth station. In this case the terrestrial tail is negligible. The second obviously occurs when the terrestrial tail is not negligible. The easiest, and most versatile design results when the tail is negligible. The case where the tail is not negligible just requires that there be terrestrial links of size indicated by Table 3.1

Satellite communication channels are naturally $(1,k)$. The large coverage of the beam allows for the " k ". In addition, with cooperation between earth stations, it is easy to change the origin of the uplink or the " 1 " in the $(1,k)$. Thus the satellite provides an easily switched, $(1,k)$ link. Generally, if the satellite provides $N(1,k);v$ links, the same satellite capacity will be used by $1(1,w);N*v$ link. This means that the Fully Connected Network using $(1,1);v$ links would use $N(N-1)*v$ units of satellite capacity while the Dynamic Allocation Network would use $N*v$ units of capacity. The Switched Broadcast Network would use v units of capacity. Furthermore a satellite network could vary the assignment of the capacity. However, those systems that require remotely induced switching (see Table 3.2) would require at least one round trip delay (.27 sec) to initiate a switch and another roundtrip delay before receiving the new view. Thus there is an unavoidable delay of .27 sec and a potential delay of .54 sec. For a satellite network, Table 3.1 indicates the number of inward and outward ports required and Tables 3.2 and 3.3 the satellite capacity. Note that a satellite network service that tariffs via ports (eg ISACOMM) will have higher costs than one with private earth stations accessing a portion of a satellite transponder.

4.3 FLEXIBILITY COMPARISONS FOR THE NETWORKS.

As mentioned earlier, there is a tradeoff possible between image quality and the degree of presence within the conferencing. The image quality (or voice quality) is captured by the bandwidth v of a connection. One presumes that a $(1,1);v$ connection will give a lower quality image than a $(1,1);2v$ connection. It is not clear at this time whether, when there are bandwidth constraints, whether it will be preferable to have more presence or better image quality. The question then to be asked is how flexible is the network a trading off image quality with presence. During a trial phase, a flexible network allows experimentation with the tradeoffs. During an operational phase, a flexible network allows graceful degradation under failure.

Although the canonical networks derived in section 4.1.1 were determined as optimal for one or more of the service networks, it is possible to size them (i.e. give the links capacity) to satisfy the service network requirements. It is assumed that if additional switching is required, in any canonical network, it will be added. The table shown below indicates the capacity that would be available to a coder in any of the configurations for each of the service networks. The networks are sized to be roughly comparable. This means that in general there are about $N(N-1)$ $(1,1);v$ simplex links total in the network. This does not mean that the route mileage will be the same for all the terrestrial networks but rather that the service networks for which they were designed will accommodate coders with bandwidth " v ". With this in mind the satellite network is sized to accommodate " v " coders for a fully connected service network. The coder bandwidth computation is based upon the worst case situation and does not use any residual capacity that the network might have for dynamically increasing the bandwidth of just a portion of the coders. Since uneven perceptual coder performance is probably a bad idea anyway, there would be little use for this feature.

Table 4.1 Comparison of Coder Capacity of The Canonical Networks

\ SERVICE ----- CANONICAL \	Fully Connected	Dynamic Allocation	Complete Broadcast	Switched Complete Broadcast	Switched Broadcast
Fully Connected	v	v	v	$(N-1)v$	$(N-1)v$
Simplex Tree (Collection)	$v/(N-1)$	v	$v/(N-1)$	v	v
Simplex Tree (Broadcast)	$v/(N-1)$	$v/(N-1)$	v	v	v
Duplex Tree	$\frac{1}{2} \frac{v}{N}$	v	v	v	$\frac{1}{2} \frac{v}{N}$
Modified Ring	$4v/(N+1)$	v	v	v	$(N-1)v$
Star (N-1)v	v	v	v	$(N-1)v$	$(N-1)v$
Satellite N(N-1)v	v	$(N-1)v$	$(N-1)v$	$(N-1)v$	$N(N-1)v$

The above table summarizes the results on the versatility of the various network alternatives. The most striking aspect of this table is that even if there were unlimited switching, the coder capacity of the terrestrial networks is in some cases reduced by a factor of N when the optimal network is used in the wrong way. This really means that any residual capacity within the network is not really usable for alternate conference spaces. The sole exception might be the graphics space because of its bursty traffic nature. The satellite network on the other hand is totally able to use its capacity and hence has the option of trading off the coder quality with presence. The second point is that except for dynamic allocation, the use of switching to increase image quality does not work. The excess capacity within the network is

just not available. These networks are in fact connectivity constrained.

The above table also makes a second very important point. The codec bandwidth capacity of a network is very dependent upon the network topology. This comment is true even if the networks were operated in a packet switching mode. The reason is that the networks must support all the simultaneous flows of the conference. There is little statistical averaging to exploit. This means that the conferencing capability of any given network is going to be strongly dependent upon its detailed structure.

5. THE INTRA CLUSTER NETWORK

As mentioned earlier, a cluster of sites is defined as those sites that can be connected via analog video. Analog video can be distributed via coax or optical fibre. A general local network architecture is shown in Figure 5:1. This shows that there is a "cluster controller" that acts as the interface to the inter cluster network. At each site there is a "site controller" that acts as the interface to the local network. All local clusters will have this architecture. However those clusters with only one site will have a very simple local network. The network shown in Figure 5:1 looks remarkably like a bus or cable distribution network. This is indeed intentional but the arguments favouring that design will be given. The major choices in the design involve

- single or multiple coaxes ... single coaxes will need multiplexing.
- video and analog codec location ... site controllers or cluster controller.
- access and transmission structure for the data, signalling, and control portions of the network. The major non control user of this part is the graphics space.

This design is predicated upon the realization of a Virtual Space. It will be adaptable to other forms of conferencing with little difficulty.

A major requirement of the intracuster network is that it be able to easily accomodate site moves and additions. This requirement favours a bus like distribution with a passive interconnect. The video is expected to be continuous so that a fixed channel allocation is desirable. Both of these imply that a multiplexed video is the best. The audio is not as continuous. It is clear that the audio must be broadcast locally and sent to the cluster controller for inter cluster broadcast. Simple coders may be used for local broadcast. The complex coders for intersite distribution may be located at the cluster controller. For the "enhanced audio space", there is only one high quality coder required per cluster. If this high quality coder is placed at the site controller, and there are K sites, there will be K instead of one encoder. The simplest solution for the audio is to multiplex the K audio channels onto a broadcast network. If it is assumed that there are K sites and that the "size" of the conference is P. The nominal P is 5. In the previous sections, $P=N$, where N is the total number of sites. This means that each site has the following:

Site Conferencing Equipment

- (P-1) cameras, monitors, and audio speakers
- 1 microphone
- 1 graphics space

If there are K local sites, these can be thought of as a mini virtual space. There K of the P cameras at any site are distributed locally and do not (necessarily) involve the cluster controller. This means that $K(K-1)$ video channels are supplied to service the local requirements only. A network of K sites has $K(P-1)$ cameras and $K(P-1)$ monitors. If there are no video codecs at the site controller, and all the channels must be routed via the cluster controller, then a total of $2K(P-1)$ video channels are needed. This will not usually be necessary for the local video channels as the camera output is compatible with the monitor input. Thus the $2K(K-1)$ channels that would have been necessary for the "mini" conference is in reality $K(K-1)$ channels. Thus the local network must provide $2K(P-1) - K(K-1) = K(2P-K-1)$ video channels. For $K=3$ and $P=5$ this is 18. Most coax systems easily carry 35 channels so that this is a feasible number. Each site controller must interface to $2(P-1)$ video channels and the cluster controller must interface to $2K(P-K)$ video channels. Thus it must have $K(P-K)$ (nominally 6) modulators and $K(P-K)$ demodulators. Thus it must also have $K(P-K)$ video codecs. If codecs were individually supplied to each of the K sites there would be $K(P-1)$ video codecs within the cluster.

Since there are exactly P audio channels, K of which originate locally and $P-K$ of which originate at the cluster controller, a simple solution is to supply a two channel system employing controlled access to the cluster controller and a queued broadcast channel from the cluster controller. This channel can also be used for the graphics space. The reason for the asymmetry is that the cluster controller will originate by far the most traffic. The total bandwidth of the queued broadcast channel is less than $P(x+g)$ where x is the bandwidth of the encoding for the voice on the local network and g is the bandwidth allocated to the graphics space. For reasonable P, and extravagant "x" (i.e 64 Kbps), the total bandwidth of this channel need be no more than 500 Kbps. This system could be thus implemented on a 1 Mbps channel with controlled access. If the data rate were kept this low, this control/graphics/audio channel could operate at baseband with little interference with the video channels. A standard Ethernet at 10 Mbps potentially interferes with analog video. If an Ethernet cable were used, then the system would have to be dual coax, one for video and one for control. Note that an unmodified Ethernet could not implement the dual channel system as its access protocol is fixed as part of the specification.

The alternative to a coax based system is direct wiring of each function. For the K,P local cluster this means that there would be

- $K(2P-K-1)$ coax cables

- P multidrop twisted pair for audio
or K+P twisted pair with a switch at the cluster controller
- P multidrop twisted pair for the graphics space
or K+P twisted pair with a switch at the cluster controller
- 2K twisted pair for control.
- Totals: $K(2P-K-1)$ coaxes
 $4K+2P$ twisted pair (assuming switch)

The resulting network will be difficult to configure and to reconfigure. In most telephone systems the solution to the multicable problem is to terminate all cables on a common patch panel. This considerably reduces the interconnect problem. In this case the total number of coaxes in the system increases to $2K(P-1)$. For $K=3$, $P=5$, this is 18 coaxes and 22 twisted pair. For sites that are exceptionally close, or where the wiring is already in place, the multiple coax may be the cost effective solution. For the general case it almost certainly will not.

There are several operational advantages of the multiplex coax approach. Some of these are as follows:

- the camera/codec, codec/monitor, or camera/monitor connections can be performed via local channel selection or via a remote signal sent by the cluster controller. This capability is important to ensure that the participant images are placed in the correct positions at all sites.
- the addition of a site requires the connection to a single cable entry point.
- there is a relatively simple change to the access scheme to allow consistent entry of a new site.
- not all sites in a local cluster need take part in a conference.
- the structure supports multiple conferences amongst mutually exclusive sets of sites.

The disadvantages of the multiplexed coax network are all related to the cost of providing the frequency converters for the channelization. It is expected that the cable runs will be sufficiently short that amplifiers will not be needed.

Thus the preferred transmission medium for the intra cluster network will employ one or two broadcast cables. The video will be frequency multiplexed and the voice, data, and control will be carried on a multiple access digital channel. Two options exist for the access control on the digital channel. The first would use a simple polling scheme initiated by the cluster controller. Since the number of sites is small, and since the conference cannot continue, at least for the cluster sites, if the cluster controller goes down, the dependence on the controller is not really an issue. This system solution would alternate between the controller broadcasting, and the local sites transmitting. The total data rate required, even for 5 or 6 sites would be less than 1 Mbps. The second solution would use a standard Ethernet protocol for the cable multiple access. This system would handle all the voice and data requirements but would require a second cable for video if it were run at the specified 10 Mbps.

6. THE SITE CONTROLLER

The (K,P) cluster is assumed to have K sites while the conference has P participants. The site controller is required to interface the site conferencing equipment to the local network. It is assumed that the complex coders, both for video and audio are located at the cluster controller. Thus the video is carried by the local network as an analog signal, and the audio is carried either analog or using simple digital codecs. In fact there may not be a need to supply any codecs for the audio. It is possible, under the appropriate circumstances, to include the audio as part of the composite signal demodulated by the TV sets. Indeed, this is the simplest way of distributing the intracluster portion of the virtual space. Thus, the site controller must supply the following:

- Cable terminator and preamp (if necessary)
for P-1 cameras (P=5 in nominal case).
- Modulator for Video plus Audio to fixed standard
Cable TV channel for P-1 cameras.
 - * Optional
 - > Channel may be locally selectable
 - > Channel may be remotely selectable
- Converter from standard Cable TV channel to fixed
VHF channel for P-1 channels.
 - * Optional
 - > Channel may be remotely selectable
- RF compatible Ethernet like interface for Graphics
Space, control, and audio (if necessary). This
equipment could be custom assembled.
- Simple control processor (to be determined)

The site controller is rather simple. The most expensive items are the video modulators. The Ethernet type access controller is the next most expensive item. The major decision for this access controller is whether to be baseband (DEC/INTEL/XEROX), or at RF (MITRENET). With most commercial offerings, there is a protocol that must also be accepted.

7. THE CLUSTER CONTROLLER

The (K,P) cluster is assumed to have K sites while the conference has P participants. The cluster controller must interface the local intracluster network with the intercluster network. It also serves as the host for the specialized audio and video codecs. As the host it must supply any logic that is required to implement any intracluster codec sharing that is desired. The interface requirements are as follows:

Intracluster Network Interface: (K=3, P=5 for nominal case)

- Modulator for Video plus Audio for each of the
K(P-K) incoming Video signals and associated
(P-K) Audio signals.
- For Virtual Space, there are
K(P-K) video encoder/decoders
K voice encoders
(P-K) voice decoders
- For Dynamic Allocation, there are
(P-K) to K(P-K) video encoders
X video decoders, $(P-K) \leq X \leq K(P-K)$
(K(K-P) - X) video frame repeaters
K "interrupt" quality voice encoders
(P-K) "interrupt" quality voice decoders
1 to K high quality voice encoders
1 to K high quality voice decoders
* local allocation logic when there
are fewer codecs than signals
- Message Forwarder for Graphics Space

Intercluster Network Interface

Terrestrial Network

- see section 8. for an example

8. A DEMONSTRATION NETWORK

A demonstration network was to have been deployed in 1982 to demonstrate the concepts of virtual space and dynamic allocation. Although this particular demonstration has been abandoned, it is still instructive to go through the design to see what the resulting network would have been.

The demonstration was to have consisted of three clusters, all in the Washington D.C. area. One cluster would have had three sites and the other two, one site apiece. The video codecs for the virtual space were to be 9.6 Kbps. The audio codecs were to be 2.4 Kbps. The intercluster network must support both virtual space (VS) and dynamic allocation for the Pseudo Virtual Space (PVS). For 3 clusters, a ring and fully connected network have the same capacity. The following table represents the flows necessary from cluster to cluster:

Table 8.1 Intercluster Flow Requirements (Demonstration Network)
Virtual Space

To From	K=3	K=1	K=1	Notes: G: Graphics Space "g" is ring g = 9.6 Kbps V: Video v = 9.6 Kbps A: Audio a = 2.4 Kbps
K=3	V: A: G:	V:3v A:3a G:g	V:3v A:3a G:g	
K=1	V:3v A:a G:g	V: A: G:	V:v A:a G:g	
K=1	V:3v A:a G:g	V:v A:a G:g	V: A: G:	

The topology for this case is given in Figure 8:1. The total bandwidth of each of the links out of the K=3 cluster is $3(a+v)+g = 45.6$ Kbps. The link connecting the two K=1 clusters has total bandwidth $(a+v+g) = 21.6$ Kbps. For the simple tree network in Figure 8.2, which also satisfies the flow constraints, the link bandwidths need to be $4(a+v)+g = 57.6$ Kbps. Thus a 56 Kbps full duplex link almost satisfies the flow requirements. For dynamic allocation, the flow requirements are given in Table 8.2.

Table 8.2 Intercluster Flow Requirements (Demonstration Network)
Pseudo Virtual Space (Dynamic Allocation)

To From	K=3	K=1	K=1	
K=3	V: A: G:	V:v A:2a+b G:g	V:v A:2a+b G:g	Notes: G: Graphics Space "g" is ring g = 9.6 Kbps V: Video v = 9.6 Kbps A: Audio a = 2.4 Kbps
K=1	V:3v A:a G:g	V: A: G:	V:v A:a G:g	
K=1	V:3v A:a G:g	V:v A:a G:g	V: A: G:	

A topology for this case is given in Figure 8.3. Note that although this is a modified ring, it has the same topology as the fully connected network. The total bandwidth of all three links is $2v+a+b+g = 40.8$ Kbps. A simple tree topology is given in Figure 8.4. Here the $K=1$ to $K=3$ flows dominate the bandwidth and the bandwidth required is $4v+b+g = 57.6$ Kbps. This means that a 56 Kbps link will almost satisfy the bandwidth requirements. Hence a network consisting of three 56 Kbps links in a triangle with the appropriate multiplexers will satisfy the flow constraints. In addition the tree of two 56 Kbps links will almost satisfy the flow constraints for both the Virtual Space and the Pseudo Virtual Space with dynamic allocation.

For each 56 Kbps modem, there must exist a multiplexor that can create an almost arbitrary, but fixed, combination of 9.6 Kbps and 2.4 Kbps channels. In addition, for the tree networks there must exist local switches to change the interconnection of some of the 9.6 Kbps channels for video and to connect the 9.6 Kbps enhanced audio channel.

Because of the switching required to implement the dynamic allocation in a tree network, the 56 Kbps RING is the RECOMMENDED topology.

The basic block diagram for the cluster controller is given in Figure 8.5. It is broken up into three modules, the Local Network Module, the Codec Module, and the Intercluster Network Module. Included in the Codec Module is the intelligence for handling the Graphic Space and site control. For single site clusters ($K=1$), the Local Network Module just interfaces the site. For the two sites with $K=1$, there is really no need for a site controller. Thus the requirements for sites One and Two are as follows:

Sites One and Two: K=1, P=5

Site Controller (Not required)

Cluster Controller

Local Network Module

- direct cable interfaces for microphone, 4 cameras, 4 TV screens
- graphics space interface

Codec Module

- 4 video codecs
- 4 audio codecs (4 interrupt quality decoders,
1 interrupt quality encoder,
1 high quality decoder,
1 high quality encoder)
 - * switch between interrupt and high quality encoders
 - * routing switch for high quality decoder output speaker
 - * output channel select for each of the video codecs.
- 1 graphics message handler
- 1 control message handler

Intercluster Module

- 2 56 Kbps full duplex modems
- 2 56 Kbps multiplexors to derive (statically) up to
4 9.6 Kbps and 4 2.4 Kbps subchannels. Subchannels on
the two modems must be interconnectable.

For Site Three, an intracluster network must be provided. Thus the requirements are as follows:

Site Three: K=3, P=5

Site Controller (3 required)

- direct cable interfaces for microphone, 4 cameras, 4 TV screens
- cable TV modulator for 4 channels (statically selectable)
- 4 cable TV converters for incoming channels
- data channel interface for graphics and control
 - * MITRENET or ETHERNET

Cluster Controller

Local Network Module

- cable TV modulator for 6 channels (statically selectable)
- cable TV converters for 6 channels
- data channel interface for graphics and control

Codec Module

- 6 video codecs
- 3 audio codecs (2 interrupt quality decoders,
3 interrupt quality encoders,
1 high quality decoder,
1 high quality encoder)
 - * switch between interrupt and high quality encoders
 - * routing switch for high quality decoder output speaker
 - * output channel select for each of the video codecs.
- 1 graphics message handler
- 1 control message handler

Intercluster Module

- 2 56 Kbps full duplex modems
- 2 56 Kbps multiplexors to derive (statically) up to
4 9.6 Kbps and 4 2.4 Kbps subchannels. Subchannels on
the two modems must be interconnectable.

It is rather obvious that the detailed interconnection of the elements is not specified above. However the general requirements should be rather obvious.

9. CONCLUSIONS

This report has discussed in some detail the requirements and constraints inherent in designing a network to support various types of teleconferencing. A range of conferencing options have been considered with the differentiations between the types primarily motivated by the networks required to transport the images, graphics, and audio. Several surprising conclusions about the inflexibility of terrestrial networks to handle a range of teleconferencing options emerge. The summary is that terrestrial networks do not adapt to different teleconferencing regimes very well. A second noticeable aspect is that although simplex links are useless for 2-point connections, they may be more economical for N-point connections. There is still more work to be done to confirm this conjecture. The topological design solutions given in this report concentrate on symmetric flow requirements for the conference and single site clusters. The general problem of multisite clusters is still unsolved.

Satellite networks on the other hand are very flexible. This is primarily due to the combined features of easily allocatable satellite bandwidth and "free" (no additional satellite bandwidth) broadcast capability. The conjecture is that satellite networks will be the transport medium of choice in future teleconferencing networks.

In addition to the general framework for teleconferencing that has been developed, a specific design problem for a specialized demonstration network has been described. The general requirements have been given. It is presumed that these are in sufficient detail for potential pricing. For the most part, the system would use readily available hardware. These specifications are given in Section 8.

10. FUTURE WORK

This report has begun a study in the design of conferencing networks that support a variable conference environment including different, audio, video, and graphics spaces. However, the study was of necessity very limited. It was primarily restricted to topological design in a dedicated very private network. Several important areas still need study. These include

- * The design of networks where there are several simultaneous conferences amongst a large number of nodes. This implies work on the following problems.
 - * A conference preparation protocol to ensure that the material needed for the participants are all in place.
 - * A connection protocol that ensures that the material needed for the conference is in place and the participants are all consistently placed. A connection should be easily invoked by a user.
 - * An integration plan for integrating the three different streams within the network. There are many options for this integration.
 - * A network topology design that accounts for switching and preconference preparation.

These are all very general problem statements and can be made more precise within the context of any apriori networking or conferencing constraints that might exist.

FIGURES

Figure 2.1:1 The Conference Model




Figure 2.1:2 The Virtual Space Teleconference




Figure 2.1:3 The Simple Teleconference






Figure 3.1:1 Two Clusters (Unlinked)



| < , > | Input Port

< | , | > Output Port

Figure 3.1:2 Two Clusters (Linked directly)



| < , > | Input Port

< | , | > Output Port

Figure 3.1:3 Two Clusters (Linked via a simplex link)

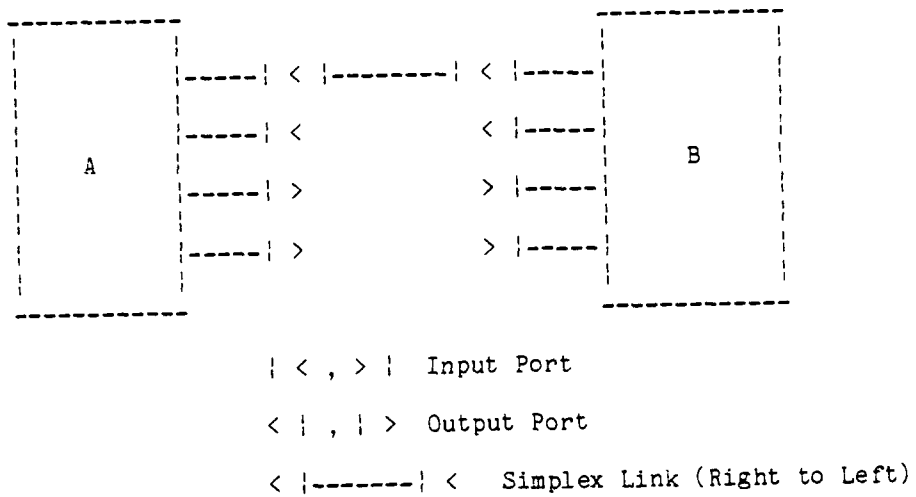


Figure 4.1:1 Example Topology of 7 Clusters Collection at A




Figure 4.1:2 Example Topology of 7 Clusters
Collection at A:Simplex Capacities




Figure 4.1:3 Example Topology of 7 Clusters
Collection at A:Duplex Capacities




Figure 4.1:4 Example Topology of 7 Clusters
Modified Ring: Duplex Capacity




Figure 4.1:5 Example Topology of 6 Clusters
Modified Ring: Duplex Capacities




Figure 4.1:6 Example Topology of 5 Clusters
Modified Ring: Routing




Figure 4.1:7 Example Topology of 6 Clusters
Modified Ring: Routing




Figure 5.1 Intracluster Network




Figure 8.1 Experimental Network
Fully Connected: Virtual Space




Figure 8.2 Experimental Network
Tree Network: Virtual Space




Figure 8.3 Experimental Network
Modified Ring:Pseudo Virtual Space




Figure 8.4 Experimental Network
Tree Network:Pseudo Virtual Space





Figure 8.5 Cluster Controller Block Diagram



DATE
ILME
—8